

CLAIMS

What is claimed is:

1 1. A computer implemented method for communicating data in a clustered
2 computing system, the method comprising:
3 detecting an occurrence of a first event at a first node of the system;
4 detecting an occurrence of a second event at the first node of the system;
5 determining that
6 the first event is identical to the second event, and
7 the occurrence of first event and the occurrence of the
8 second event are therefore a first occurrence and a
9 second occurrence of an identical event;
10 appending a notification including the information describing the
11 identical event onto an existing message, wherein
12 the notification is appended only once for both the first
13 occurrence and the second occurrence,
14 the notification is not appended twice, once for
15 each occurrence, and
16 the message was destined to be propagated to a
17 second node; and
18 propagating the message along with the information describing the
19 identical event to the second node.

1 2. The method of claim 1 wherein the message was generated for purposes
2 other than sending information appended.

1 3. The method of Claim 1, wherein the determining further comprises

2 comparing information that describes the first event with information
3 that describes the second event to determine whether the first
4 event and the second event are identical; and
5 the method further comprising if the two events are identical, then indicating
6 that the information that describes the second event no longer needs to
7 be retained.

1 4. The method of Claim 1 wherein the method further comprises:
2 setting an identifier indicating that the information
3 describing the identical event is to be appended
4 onto a message and propagated to a particular
5 node.

1 5. The method of Claim 1, wherein said clustered computing system
2 comprises a database management system.

1 6. The computing environment of Claim 1, wherein said clustered
2 computing system comprises a shared-disk database system.

1 7. The computing environment of Claim 1, wherein said clustered
2 computing system comprises a shared-cache parallel database management
3 system.

1 8. The computing environment of Claim 1, wherein said clustered
2 computing system comprises a shared-nothing database management system.

1 9. The computing environment of Claim 1, wherein said clustered
2 computing system comprises a distributed database management system.

- 1 10. The method of Claim 1, wherein the method further comprises:
2 searching a shared-memory event buffer having a size that is fixed.
- 1 11. The method of Claim 1 wherein the message has a fixed size, and the method further
2 comprises:
3 appending additional information that describes additional events onto existing
4 message traffic until free space in the fixed-size message is filled.
- 1 12. The method of Claim 1, wherein the method further comprises
2 placing the information describing the identical event in a queue.
- 1 13. The method of claim 12, wherein the queue includes at least a priority queuing
2 mechanism in order to determine a priority for events such that high priority events would
3 supercede a low priority events in an event notification queue.
- 1 14. The method of claim 1 wherein an in-memory hash index is used to determine if an
2 event exists in a shared-memory event buffer.
- 1 15. The method of claim 1 wherein the method further comprises:
2 partitioning a shared-memory event buffer;
3 generating an event buffer entry of the shared memory event buffer;
4 placing an event identifier into the event buffer entry; and
5 inserting the information describing the identical event into the event buffer
6 entry.
- 1 16. The method of claim 15 the method further comprises if between a fastest head
2 pointer and a tail pointer there does not exists a buffer entry in the shared memory event
3 buffer for the identical event, generating a new event buffer entry, and the inserting further

4 comprises inserting the information describing the identical event into the new event buffer
5 entry.

1 17. The method of claim 15 wherein the inserting comprises:
2 if between a fastest head pointer and a tail pointer there exists a buffer entry in the shared
3 memory event buffer for the identical event, updating the buffer entry so that the buffer entry
4 represents the second occurrence.

1 18. The method of claim 15 further comprising
2 using a round robin method and the shared memory event buffer to determine
3 to which existing message to appended the information describing the
4 identical event.

1 19. The method of Claim 1, wherein the method further comprises the step
2 of:
3 maintaining the information that describes a plurality of events.

1 20. The method of Claim 19, wherein the method further comprises the step
2 of :
3 maintaining the information that describes the plurality of events
4 in a shared-memory event buffer .

1 21. The method of Claim 19, wherein the method further comprises:
2 maintaining the information that describes the plurality of events
3 in a circular buffer.

1 22. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 1.

1 23. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 2.

1 24. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 3.

1 25. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 4.

1 26. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 5.

1 27. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 6.

1 28. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 7.

1 29. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 8.

1 30. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 9.

1 31. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 10.

1 32. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 11.

1 33. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 12.

1 34. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 13.

1 35. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 14.

1 36. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 15.

1 37. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 16.

1 38. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 17.

1 39. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 18.

1 40. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 19.

1 41. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 20.

1 42. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 21.

1 43. A computer implemented method for communicating data in a clustered
2 computing system, the method comprising:
3 detecting a generation of an event at a first node of the
4 system,

5 determining if the event is identical to another event that
6 was previously generated,
7 appending only one copy of the information is appended
8 onto an existing message whether or not the other
9 event is identical, wherein the message was
10 destined to be propagated to a second node; and
11 propagating the message along with the information
12 describing the identical event to the second node.

1 44. The method of claim 43, wherein the determining further comprises:
2 if
3 there exists a stored indication that an identical event was previously
4 generated and
5 the propagating of the message having the information appended did not
6 yet occur,
7 then an indication is stored that multiple identical events were generated.

1 45. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 43

1 46. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 44.

1 47. A computer implemented method comprising the computer-
2 implemented steps of:
3 receiving, at a database server that is executing on a second node
4 in a distributed system, a message that was transmitted by
5 a first node in the distributed system wherein

6 the message has appended thereon information that
7 describes an event that occurred at least a
8 second time before the message was sent, and
9 a second description of the event is not received
10 despite the second occurrence; and
11 retrieving the information that describes the event signaling processes in the
12 local node that said event has occurred.

1 48. The method of claim 47, wherein the method further comprises:
2 invoking concurrency control techniques to control concurrent access to a
3 shared-memory event buffer from processes that propagate messages to
4 subscriber nodes and processes that generate events.

1 49. The method of claim 47, wherein said method further comprises:
2 maintaining information that describes a plurality of events,
3 coalescing the information that describes a plurality of events, wherein the
4 information that describes a plurality of events may be coalesced for
5 the same event into a single event notification.

1 50. The method of Claim 47, wherein the method further comprises:
2 maintaining the information that describes a plurality of
3 events.

1 51. The method of Claim 47, wherein the method further comprises step of:
2 maintaining the information that describes the plurality of
3 events in a shared-memory event buffer .

1 52. The method of Claim 47, wherein the method further comprises the step
2 of:

3 maintaining the information that describes the plurality of
4 events in a circular buffer.

1 53. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 47.

1 54. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 48.

1 55. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 49.

1 56. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 50.

1 57. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 51.

1 58. A computer-readable medium carrying one or more sequences of
2 instructions, which when executed by one or more processors, causes the one or
3 more processors to perform the method recited in Claim 52.

1 59. A computer apparatus comprising:
2 a processor; and

. . .

3 a memory coupled to the processor, the memory containing one or more sequences of
4 instructions for event notification in a clustered computing environment,
5 wherein execution of the one or more sequences of instructions by the
6 processor causes the processor to perform the method of claim 1.